ANS MC2015 - Joint International Conference on Mathematics and Computation (M&C), Supercomputing in Nuclear Applications (SNA) and the Monte Carlo (MC) Method · Nashville, Tennessee · April 19–23, 2015, on CD-ROM, American Nuclear Society, LaGrange Park, IL (2015)

arXiv:1702.02111v1 [cs.CE] 7 Feb 2017

RAYLEIGH QUOTIENT ITERATION WITH A MULTIGRID IN ENERGY PRECONDITIONER FOR MASSIVELY PARALLEL NEUTRON TRANSPORT

R.N. Slaybaugh[1], **T.M. Evans**[2], **G.G. Davidson**[2], **and P.P.H. Wilson**[3]
[1] Department of Nuclear Engineering, University of California, Berkeley, 4173 Etcheverry Hall, Berkeley, CA 94720, slaybaugh@berkeley.edu
[2]Radiation Transport Group, Oak Ridge National Laboratory, P.O. Box 2008, Oak Ridge, TN 37831, evanstm@ornl.gov; davidsongg@ornl.gov
[3] Department of Nuclear Engineering and Engineering Physics, University of Wisconsin – Madison, 419 ERB, 1500 Engineering Drive, Madison, WI 52706, wilsonp@cae.wisc.edu

ABSTRACT

Three complementary methods have been implemented in the code Denovo that accelerate neutral particle transport calculations with methods that use leadership-class computers fully and effectively: a multigroup block (MG) Krylov solver, a Rayleigh quotient iteration (RQI) eigenvalue solver, and a multigrid in energy preconditioner. The multigroup Krylov solver converges more quickly than Gauss Seidel and enables energy decomposition such that Denovo can scale to hundreds of thousands of cores. The new multigrid in energy preconditioner reduces iteration count for many problem types and takes advantage of the new energy decomposition such that it can scale efficiently. These two tools are useful on their own, but together they enable the RQI eigenvalue solver to work. Each individual method has been described before, but this is the first time they have been demonstrated to work together effectively.

RQI should converge in fewer iterations than power iteration (PI) for large and challenging problems. RQI creates shifted systems that would not be tractable without the MG Krylov solver. It also creates ill-conditioned matrices that cannot converge without the multigrid in energy preconditioner. Using these methods together, RQI converged in fewer iterations and in less time than all PI calculations for a full pressurized water reactor core. It also scaled reasonably well out to 275,968 cores.

Key Words: Krylov, Rayleigh Quotient, Multigrid, Preconditioning

1 INTRODUCTION

The goal of this research is to accelerate neutral particle transport calculations with methods that use large computers fully and effectively, facilitating the design of better nuclear systems. Typical steady-state deterministic transport problems today are three-dimensional, have up to thousands $\times$ thousands $\times$ thousands of mesh points, use up to $\sim$150 energy groups, include accurate expansions of scattering terms, and are solved over many angular directions. The next generation of challenging problems are even more highly refined and the flux must be calculated quickly and accurately. Very large computers, like Titan [1], are now available to perform such high-fidelity calculations. New solution methods must outperform existing ones that are not able to take full

advantage of new computer architectures or have convergence properties that limit their usefulness for difficult problems.

Three complementary methods have been implemented in the code Denovo [2] that accomplish this goal: a multigroup block (MG) Krylov solver, a Rayleigh quotient iteration (RQI) eigenvalue solver, and a multigrid in energy (MGE) preconditioner. Each individual method has been described before (see [3], [4], [5]), but this is the first time they have been demonstrated to work together effectively.

The MG Krylov solver was designed to improve performance when compared to Gauss Seidel iteration (GS) and to dramatically increase the number cores Denovo can use. Instead of sequentially solving each group with some inner iteration method and then using GS for outer iterations to converge the upscattering (which can be time consuming when many groups contain upscattering), the MG Krylov solver treats a block of groups at once such that the inner-outer iteration structure is removed. This results in faster solutions for most problem types. In addition, the block Krylov solver allows energy groups to be solved simultaneously because the multigroup-sized matrix-vector multiply can be divided up in energy and parallelized–extending the number of cores that can be used efficiently by Denovo from tens of thousands to hundreds of thousands.

An MGE preconditioner was added to Denovo to reduce iteration count for all problem types and to address convergence issues associated with RQI. The preconditioner uses a multigrid method in the energy dimension. A set of energy grids with increasingly coarse energy group structures are created. This is implemented in a way that easily and efficiently takes advantage of the new energy decomposition. The multigrid algorithm is applied within each energy set such that the energy groups are only restricted and prolonged between groups on that set. Inter-set communication in the preconditioner is kept low, so the scaling in energy is very good.

Theory indicates that RQI should converge in fewer iterations than traditional eigenvalue solvers like power iteration (PI), particularly for problems that are challenging for those solvers. However, RQI causes the systems it operates on to become nearly singular, and thus, it requires a preconditioner. Further, the implementation of RQI results in a set of equations that is mathematically equivalent to having upscattering in every group, so the scattering matrix becomes energy-block dense. Handling energy-block dense systems when there are more than a few energy groups is not tractable with GS as the multigroup solver. It is only the MG Krylov solver that makes RQI reasonable to use when there are many energy groups.

The remainder of this paper details why these methods are so complementary and presents results demonstrating this behavior. Section 2 discusses each of the new methods in the context of commonly-used methods. Section 3 gives an overview of relevant past work. New results from using the three new methods together are shown in Section 4, and concluding remarks are made in Section 5.

2 BACKGROUND

The eigenvalue form of the multigroup S_N equations can be written in operator form as

$$\mathbf{L}\psi = \mathbf{MS}\phi + \frac{1}{k}\mathbf{M}\chi\mathbf{f}^T\phi\,. \tag{1}$$

Here $\mathbf{L}$ is the first-order linear differential transport operator; $\mathbf{M}$ is the moment-to-discrete operator that projects the angular flux moments, ϕ, onto discrete angles; $\mathbf{S}$ is the scattering matrix; $\mathbf{f}$ contains the fission source, $\nu\Sigma_f$; and k is the asymptotic ratio of the number of neutrons in one generation to the number in the next. This can be converted to a fixed-source equation by replacing the fission term with an external source, q. The angular flux moments are related to the angular flux through the discrete-to-moment operator: $\phi = \mathbf{D}\psi$. Using this relationship, Eq. (1) can be rearranged such that it is a function of only ϕ. The formulation is aided by defining $\mathbf{T} = \mathbf{DL}^{-1}$ and $\mathbf{F} = \chi\mathbf{f}^T$ [6]:

$$(\mathbf{I} - \mathbf{TMS})\phi = \frac{1}{k}\mathbf{TMF}\phi\,. \tag{2}$$

Once the matrices are multiplied together, a series of single "within-group" equations that are each only a function of space and angle result. If the groups are coupled together by neutrons scattering from a low energy group to a higher energy group, then iterative "multigroup" solves over the coupled portion of the energy range may be required. If the eigenvalue is desired, an additional "eigenvalue" solve is needed.

2.1 Block Krylov Solver

Traditionally, the multigroup solve has been done with Gauss Seidel iteration. A space-angle solve using a within-group solver, such as source iteration or a Krylov method, is performed for each energy group in series. The groups are solved from $g = 0$, the highest energy, to $g = G$, the lowest. For a group g and an energy iteration index j this is [6]

$$\big(\mathbf{I} - \mathbf{TMS}_{gg}\big)\phi_g^{j+1} = \mathbf{TM}\big(\sum_{g'=0}^{g-1}\mathbf{S}_{gg'}\phi_{g'}^{j+1} + \sum_{g'=g+1}^{G}\mathbf{S}_{gg'}\phi_{g'}^{j} + q_g\big)\,. \tag{3}$$

The first term on the right includes downscattering contributions from higher energies, and the second term represents upscattering contributions from lower energy groups that have not yet been converged for this energy iteration. Convergence of GS is governed by the spectral radius of the system, so the method can be very slow when upscattering has a large influence on the solution [7]. GS is fundamentally serial in energy because of how the group-to-group coupling is treated.

The MG Krylov solver removes the traditional within-group, multigroup iteration structure to make one space-angle-energy iteration level. This allows the solver to handle upscattering efficiently since Krylov methods generally converge more quickly than GS for problems with upscattering, and enables parallelization in the energy dimension. The solver has been shown to successfully scale to hundreds of thousands of cores because of the addition of energy parallelization [3, 8].

The multigroup Krylov method applied to a block of groups is shown here, where $\mathbf{S}_{\text{block}}$ contains the block of upscattering groups and $\mathbf{S}_{\text{block_source}}$ has the downscattering-only groups:

$$\underbrace{(\mathbf{I} - \mathbf{TMS}_{\text{block}})}_{\tilde{\mathbf{A}}} \phi^{n+1}_{\text{block}} = \mathbf{TM}(\mathbf{S}_{\text{block_source}} \phi^{n+1}_{\text{block_source}} + q) \, . \tag{4}$$

The AztecOO package of Trilinos [9] provides Denovo's Krylov solver, with a choice of either GMRES(m) or BiCGSTAB. The Krylov solver is given an operator that implements the action of $\tilde{\mathbf{A}}$, or the matrix-vector multiply and sweep. In the MG Krylov solver, $\tilde{\mathbf{A}}$ is applied to an iteration vector, v, containing the entire block of groups instead of just one group.

To implement the energy parallelization, the problem is divided into energy sets, with groups distributed evenly among sets. After each set performs its part of the matrix-vector multiply, a global reduce-plus-scatter is the only required inter-set communication. The added energy decomposition offers the ability to further decompose a problem, even if the performance limit of spatial decomposition has been reached. The total number of cores is equal to the number of computational domains, that is, the product of the number of energy sets and the number of spatial blocks. For 20,000 spatial blocks and 10 energy sets, which is a reasonable decomposition, 200,000 cores will be used. See Ref. [3], [8], or [6] for more details. It is worth noting that the scaling limits seen in Denovo are not fundamental; good scaling beyond tens of thousands of cores is possible without energy decomposition (see, e.g., [10]).

2.2 Eigenvalue Solvers

A common way to solve k-eigenvalue problems is with power iteration. This method is attractive because it only requires matrix-vector products and two vectors of storage space.

$$\mathbf{A}\phi = k\phi \, ; \qquad \mathbf{A} = (\mathbf{I} - \mathbf{TMS})^{-1}\mathbf{TMF} \, , \tag{5}$$

$$\phi^{i+1} = \frac{1}{k^i}\mathbf{A}\phi^i \, ; \qquad k^{i+1} = k^i \frac{\mathbf{f}^T\phi^{i+1}}{\mathbf{f}^T\phi^i} \, . \tag{6}$$

PI uses the form of the problem seen in Eq. (5) and then iterates as shown in Eq. (6), where i is the iteration index. Inside of PI, the application of $\mathbf{A}$ to ϕ requires the solution of a multigroup problem that looks like a fixed source problem: $(\mathbf{I} - \mathbf{TMS})y^i = \mathbf{TMF}\phi^i$. PI's convergence can be very slow for problems of interest. For an $n \times n$ matrix $\mathbf{A}$, an eigenvalue-vector pair satisfies $\mathbf{A}x_i = \lambda_i x_i$ for $i = 1, ..., n$. Let $\sigma(\mathbf{A}) \equiv \{\lambda \in \mathbb{C} : rank(\mathbf{A} - \lambda\mathbf{I}) < n\}$ be the spectrum of $\mathbf{A}$ and the eigenvalues be ordered as $|\lambda_1| > |\lambda_2| \geq \cdots \geq |\lambda_n| \geq 0$. The error from PI is reduced in each iteration by a factor of $\mathbf{A}$'s dominance ratio, λ_2/λ_1. PI will converge slowly for loosely coupled systems because λ_2 is close to λ_1 [11].

Shifted inverse iteration (SII) typically converges more quickly than PI. SII capitalizes on the fact that for some shift μ and some matrix $\mathbf{P}$, $(\mathbf{P} - \mu\mathbf{I})$ will have the same eigenvectors as $\mathbf{P}$. If $\mu \notin \sigma(\mathbf{P})$, then $(\mathbf{P} - \mu\mathbf{I})$ is invertible and $\sigma([\mathbf{P} - \mu\mathbf{I}]^{-1}) = \{1/(\lambda - \mu) : \lambda \in \sigma(\mathbf{P})\}$. Eigenvalues of $\mathbf{P}$ that are near the shift will be transformed to extremal eigenvalues that are well separated from the others. The shifted and inverted matrix is used in a power iteration-type scheme. Given

a good shift, $\mu \approx \lambda_1$, SII usually converges more quickly than PI, especially for loosely coupled systems [11].

Rayleigh Quotient Iteration is an SII method that uses an optimal shift: the Rayleigh quotient (RQ). For a generalized eigenvalue problem $\beta \mathbf{A}x = \alpha \mathbf{B}x$, the RQ is

$$\rho = \frac{y^T \mathbf{A} x}{y^T \mathbf{B} x} \,. \tag{7}$$

If x and y are right and left eigenvectors corresponding to the scalars α and β, respectively, then $\alpha = y^T \mathbf{A} x$ and $\beta = y^T \mathbf{B} x$. In this case the system's eigenvalue is $\gamma = \frac{\alpha}{\beta}$ and $\rho = \gamma$ [12]. The RQ provides the minimum residual for the eigenvalue problem in the least squares sense. The Rayleigh quotient is thus an optimal guess for an eigenvalue given a vector close to the corresponding eigenvector. The idea of RQI is to use the RQ as the shift in SII, but the shift is updated with the new eigenvector estimate at every iteration i: $\mu_i = \gamma_{i-1} - \rho_i$. RQI has better convergence properties than SII since the RQ is an optimal guess for the eigenvalue of interest. For more details on the Rayleigh quotient and RQI in general, refer to Ref. [13].

RQI has been implemented in Denovo, as detailed in Ref. [4], by subtracting $\rho \mathbf{TMF}$ from both sides of Eq. (2). This gives the following shifted system, where $\gamma \equiv \frac{1}{k}$:

$$(\mathbf{I} - \mathbf{TM}\tilde{\mathbf{S}})\phi = (\gamma - \rho)\mathbf{TMF}\phi \,, \tag{8}$$

and $\tilde{\mathbf{S}} \equiv \mathbf{S} + \rho \mathbf{F}$. This new matrix, $\tilde{\mathbf{S}}$, is energy-block dense since the fission matrix is dense and looks like one big upscattering block. Traditional solution methods for the fixed source part of the equation do not handle dense scattering matrices well. The RQI solver added to Denovo uses the MG Krylov solver, designed to handle dense scattering matrices effectively, to overcome the use of a dense $\tilde{\mathbf{S}}$. In addition, RQI can be decomposed in energy and takes advantage of the scaling properties of the multigroup solver.

Using RQI in combination with a Krylov method, however, raises some concerns about whether or not it can converge because the matrix becomes so ill-conditioned. Peters and Wilkinson [14], however, proved that ill-conditioning is not a fundamental problem for inverse iteration methods. Trefethen and Bau [11] assert that this is the case as long as the fixed source portion is solved with a backwards stable algorithm. Paige et al. [15] demonstrated that GMRES is backwards stable when finding x in $\mathbf{A}x = b$ for a "sufficiently nonsingular $\mathbf{A}$" and define associated criteria. For ill-conditioned systems, then, Krylov methods may not be backwards stable and will tend to converge very slowly. Many researchers have found that Krylov methods must be preconditioned to achieve good convergence in practice, e.g. [11], [15]. The past work section below demonstrates that RQI *does* need preconditioning for convergence in problems of interest.

2.3 Multigrid in Energy Preconditioner

Preconditioning is needed to make the RQI algorithm tractable and is important for increasing the robustness of Krylov methods and decreasing Krylov iteration count. This is particularly true

for the multigroup Krylov solver; it can create large subspaces because it forms the subspaces with multiple-group-sized vectors. As a result, any reduction in iteration count will be of significant benefit in terms of memory and cost per iteration. A right preconditioner that does multigrid in the energy dimension, meaning the grids are in energy such that the energy group structure is coarsened and each lower grid has fewer groups, and is designed to work with the MG Krylov solver was implemented in Denovo [5]. To understand why multigrid in energy makes sense for neutron transport, some highlights about these methods are discussed here (see Ref. [16] for details).

The error in x_i, the ith guess for $\mathbf{A}x_i = b_i$, can be written as a combination of Fourier modes. Each Fourier mode has a frequency; low-frequency modes are smooth and high-frequency modes are oscillatory. Stationary iterative methods remove high-frequency error components quickly but can take many iterations to remove the low-frequency ones. Multigrid methods use the smoothing effects of iterative methods by making smooth errors look oscillatory and thus easier to remove. These methods remove the low-frequency error modes by restricting error to coarser grids, removing now higher-frequency error, and prolonging the smoothed error back to the fine grid.

An important principle is that the preconditioner is only attempting to approximately invert $\mathbf{A}$. It is therefore reasonable to use a less accurate angular discretization in the preconditioner than the rest of the code. For example, the whole problem may be solved at S_{10}, but the preconditioner could use S_2.

The user chooses the number of V-cycles done for each preconditioner application. One V-cycle proceeds from the finest grid to the coarsest grid and back to the finest. Each additional V-cycle should remove more error, but has a computational cost. The depth of the V-cycle can also be specified by the user. The default behavior is determined by the number of groups, such that the grids will be coarsened until there is only one energy group. The number of grids needed in that case is floor$\big(\log_2(G-1)\big)+2$ [5].

Some number of relaxations are performed on each level while traversing down and up the grids in a V-cycle. Performing more relaxations per grid should remove more error, but at a computational cost. The implemented relaxation method is weighted Richardson iteration with a weight, ω, that can be set by the user and defaults to 1. When applied to the transport equation, this is

$$\phi^m = \big(\mathbf{I} + \omega(\mathbf{TMS} - \mathbf{I})\big)\phi^{m-1} + \omega b^{m-1} \,. \tag{9}$$

The MGE preconditioner was designed to take advantage of the energy decomposition used by the MG Krylov method. When using multiple energy sets, each set does its own "mini" V-cycle with the groups on that set. Each set restricts, prolongs, and relaxes only its own groups and does not need to communicate with other energy sets beyond what is needed for handling upscattering in the solver. This strategy is preferable to doing a full V-cycle across sets, which would require much more data transfer and accounting. This model is also a communication savings compared to using grids in space or angle. An additional benefit is the simplicity of energy grids. Energy is one-dimensional, which allows for simpler coarsening and refinement than spatial or angular grids.

3 PAST WORK

This section summarizes results from using RQI *without* preconditioning [4] and the MGE preconditioner used with fixed source problems or PI [5]. The purpose of this section is to highlight the capabilities that have been demonstrated, and point out the short-comings that can be overcome by using the MG Krylov solver, RQI eigenvalue solver, and MGE preconditioner together. The goal of using this system of solvers is to improve convergence behavior of the multigroup Krylov solves integrated over eigenvalue solves. The best metric for measuring this is the total number of multigroup Krylov iterations used in a calculation because it encompasses the total work done. Timing comparisons should be considered heuristically unless otherwise specified.

3.1 Unpreconditioned RQI

The goal of the RQI studies that have been published was to find if RQI is useful without preconditioning. We solved several problems, where the first set was two small eigenvalue test problems, one with vacuum boundaries and one with reflecting, that had small dominance ratios. We found that RQI got the correct answer and converged in fewer iterations than PI.

However, an intermediately-sized problem did not work so well: the eigenvector did not converge after the first iteration. The value of k oscillated between 0.3966 and 0.3967 (the correct value was 0.4) until the calculation was manually terminated. In the 2-D and 3-D C5G7 MOX Benchmark problems ([17], [18]), which are more realistic calculations, RQI did not converge the eigenvector nor find an eigenvalue close to the correct one. These cases of RQI's non-convergence are caused by the non-convergence of the inner linear solve. Given a large enough subspace and/or enough iterations (which is not feasible in practice because of memory limitations), the linear solver would always converge and so would RQI.

The more challenging problems showed that, as expected, the Krylov solver often cannot converge the eigenvector with the ill-conditioned systems created by RQI in a reasonable amount of time. Work by Hamilton [19] also demonstrated the need for preconditioning in solving shifted transport problems.The small problems, however, showed that RQI can require fewer Krylov iterations than PI if the multigroup iterations are converged. Thus, if the MG Krylov solver is preconditioned so that the eigenvector converges, RQI may be able to find the correct eigenvalue more efficiently than PI for cases of interest.

This leads to the question: what preconditioner should we choose? Iterative methods reduce oscillatory but not smooth error modes effectively and the smooth error can prevent iterative methods from converging. This behavior is characterized by rapid error reduction in the first several iterations followed by very little error reduction. Such a trend was observed in the tests where RQI failed. Multigrid methods selectively remove smoother error components and are therefore ideal for accelerating this type of problem.

3.2 MGE Preconditioner

Much of the previously-published work for the MGE preconditioner focused on choosing all of the options that control the preconditioner: Richardson iteration weight, number of V-cycles per preconditioner application, number of relaxations per level, quadrature in the preconditioner, effects of energy sets, and depth of the V-cycle (Hamilton's work [19] also discusses multigrid-cycle parameter selection). The syntax used throughout this work will be that $w\#$ is the weight, $r\#$ is the number of relaxations per level, and $v\#$ is the number of V-cycles, e.g. $w1r1v1$ is one relaxation per level, one V-cycle, and a weight of 1.0. Using more preconditioning means using larger values of w and/or r and/or v.

Some initial tests provided a basis for choosing each of these. As a result, the default parameters are $w1r2v2$. Tests also showed that using a reduced quadrature set inside the preconditioner is very valuable. Another important area of investigation demonstrated that the preconditioner scaled very well with multiple energy sets. This was largely because increasing sets reduces the cost of the preconditioner since each set uses a shallower V-cycle and therefore does less work. The results of several tests additionally demonstrated that, in general, it is better to use only a few grids (shallower V-cycle).

All of these tests inform the best way to use the MGE preconditioner, but do not determine whether and when it is useful. To begin investigating this question, two eigenvalue problems were solved with preconditioned and unpreconditioned PI. The first problem was the 3-D C5G7 MOX Benchmark problem. In this case the use of MGE significantly reduced Krylov iteration count, but increased the overall runtime. A full PWR problem (described in subsection 4.3) exhibited similar behavior: fewer iterations in more time.

These results suggest that the MGE preconditioner was a failed experiment after all. However, the two eigenvalue problems solved were only solved with PI and the mathematical properties of the MGE preconditioner suggest it would benefit the RQI solver. Thus, the work published so far has not settled the question of whether MGE is a useful preconditioner for at least some problems.

4 RESULTS

The collection of observations in the past work section led to the questions (1) Will preconditioning with MGE facilitate the use of RQI? and (2) Will the combination of RQI, MGE, and the block Krylov solver be advantageous for at least some problems? This section is designed to answer these two questions. In this Section, reducing the Krylov iteration count and reducing total compute time are measures of success. Unless otherwise noted, all test problems used a step characteristic spatial solver, level-symmetric angular quadrature, and the grid depth was determined using the default approach. The Krylov solver was GMRES(m), which is set to limit the number of multigroup iterations to 1,000 if the problem does not converge earlier. The convergence tolerances are noted for each problem. The tolerance for the multigroup solve is the convergence tolerance used by GMRES in Trilinos. The eigenvalue tolerance is used by PI or RQI to determine if the eigenvalue has converged.

4.1 2-D C5G7

Using guidance from the past work, MGE was applied to the 2-D C5G7 benchmark using both PI and RQI. The goals were to see if preconditioned RQI could converge the eigenvector (flux) and solve for k, and to investigate the effect of preconditioning with both RQI and PI. The calculation used 16 cores on the small Orthanc cluster at Oak Ridge: 4 x-blocks, 4 y-blocks, 1 z-block and 1 energy set. The total and upscattering tolerances were 1×10^{-3}, with a k tolerance of 1×10^{-5}.

The first study used the MGE preconditioner with PI. The results are shown in Table I, where "Krylov" is the total number of Krylov iterations and "PI" is the total number of eigenvalue iterations. All calculated ks were within the uncertainty of the benchmark and so are not reported. This study shows the preconditioner is very effective at reducing the number of Krylov iterations used by PI, but PI is much faster without preconditioning. The unpreconditioned case, corresponding to a weight of 0, took about twice the MG Krylov iterations as the $w\#r1v1$ cases. Adding the preconditioner dramatically reduced the number of Krylov iterations but more than doubled the runtime. With $r1v1$, increasing the weight from 1 to 1.4 decreased the number of Krylov iterations and the time to solution. The time and iteration count both increased with a weight of 1.5. When the weight was increased beyond 1.5 none of the multigroup iterations converged so they are not reported here. The two calculations with a higher level of preconditioning converged in the fewest Krylov iterations, with the $w1.4r2v2$ taking the smallest amount of time and $w1r3v3$ taking the second largest.

Table I. 2-D C5G7 benchmark: preconditioner weight variation with PI

Weight	Relaxations	V-cycles	Krylov	PI	Time (s)
0.0	0	0	3,129	32	8.54×10^3
1.0	1	1	1,649	31	2.12×10^4
1.1	1	1	1,591	31	2.06×10^4
1.2	1	1	1,546	31	1.98×10^4
1.3	1	1	1,492	31	1.92×10^4
1.4	1	1	1,458	31	1.91×10^4
1.5	1	1	1,771	31	2.29×10^4
1.4	2	2	438	31	1.77×10^4
1.0	3	3	253	31	2.28×10^4

The results from the preconditioned RQI study are in Table II. In all cases, except the unpreconditioned one, k was within the uncertainty of the benchmark value. The "$<$ 1,000?" column indicates whether or not the multigroup iterations converged during the RQI process, where a number indicates the last eigenvalue iteration for which the Krylov method took less than 1,000 iterations. The relative time is the ratio of the case of interest to the unpreconditioned PI time of 8.54×10^3 seconds.

These results show a few important things. Most significantly, with enough preconditioning the multigroup iterations within RQI can be converged and the correct eigenpair can be found.

Table II. 2-D C5G7 benchmark: convergence study with RQI

Weight	Relaxations	V-cycles	Krylov	RQI	$<$ 1,000?	Rel. Time
0	0	0	119,006	120*	no	10.98
1	1	1	16,007	17	no	23.65
1.2	1	1	40,008	41*	no	13.00
1	3	1	n/a	n/a*	7	n/a
1	2	2	11,158	19	alternated	46.72
1	3	2	3,320	19	14	19.23
1	3	3	299	19	yes	3.01
1.1	3	3	281	19	yes	2.80
1.3	3	3	254	19	yes	2.57
1.5	3	3	n/a	n/a*	no	n/a

*terminated manually

This was true even when the eigenvector did not converge inside each eigen iteration (though these cases were significantly more time consuming). Further, as the preconditioning increased the eigenvector came closer to converging for all iterations. For the first three $w\#r3v3$ cases, all of the Krylov iterations converged. This test case was the first to demonstrate that the preconditioner can get RQI to converge when it did not converge without preconditioning. When the Krylov iterations did converge, increasing the weight decreased iteration count and wall time for small weights. Finally, too much weight prevented the calculation from converging.

A big question is whether RQI is faster than PI with preconditioning, but only the $w1r3v3$ calculation overlapped between RQI and PI. PI took fewer Krylov iterations, 253 compared to 299, and less time, 2.28×10^4 compared to 2.57×10^4 seconds. From the standpoint of comparing eigenvalue solution methods, it is worth noting that RQI required 19 eigenvalue iterations while PI required 31. Thus, when given eigenvectors that have been converged to the same tolerance, RQI needed fewer eigenvalue iterations than PI. For this test preconditioned RQI did not perform as well as preconditioned PI, though the times and iteration counts were close to one another.

4.2 3-D C5G7

Next, the preconditioner using both PI and RQI was applied to the 3-D C5G7 benchmark. The goals of this study were essentially the same as the 2-D study, except that this problem is larger and represents a more realistic problem. The medium-sized OIC cluster at Oak Ridge was used, and each problem was given 720 cores with 40 x-blocks, 18 y-blocks, and 5 z-blocks. The total and upscattering tolerances were 1×10^{-4}, with a k tolerance of 1×10^{-5} unless otherwise indicated. The wall time limit was 12 hours.

The preconditioned PI results were reported in [5], so we only add the RQI results shown in Table III here. The relative time is compared to unpreconditioned PI, 4.46×10^3 seconds. Cases with no or a small amount of preconditioning ($w1r1v1$, $w1.5r1v1$, $w1.2r2v1$) are not reported since none converged within the walltime limit. This indicates a small amount of preconditioning

was insufficient. However, with a lot of preconditioning ($w1r4v4$, $w1.5r5v5$) too much time was spent in the preconditioner and the problem did not finish in time either. Convergence was obtained in several cases by lowering the tolerances.

Table III. 3-D C5G7 benchmark: preconditioner parameter scoping with RQI

Weight	Relaxations	V-cycles	Krylov	RQI	Rel. Time
1.3	2	2	302	19	5.20
1	3	3	103	9*	6.67
1	3	3	164	15+	7.59
1.5	3	3	187	19	7.26
1	4	4	n/a	n/a	exceeded wall time
1	4	4	74	9*	5.13
1.5	5	5	n/a	n/a	exceeded wall time

*tol and upscatter tol = 1×10^{-5}, k tol = 1×10^{-3}
+tol and upscatter tol = 1×10^{-4}, k tol = 5×10^{-5}

With an intermediate amount of preconditioning, RQI converged and performed better than the analogous PI cases. There are three cases where both problems finished and the same tolerances were used: $w1.3r2v2$, $w1r3v3$, $w1.5r3v3$. These results are shown together in Table IV. This table displays time instead of relative time to facilitate comparison between two cases rather than across all cases. In all three pairs, the RQI calculations took less time and fewer eigenvalue iterations.

Table IV. 3-D C5G7 benchmark: RQI and PI comparison

Sovler	Weight	Relaxations	V-cycles	Krylov	Eigenvalue	Time (s)
RQI	1.3	2	2	302	19	2.32×10^4
PI	1.3	2	2	288	32	2.84×10^4
RQI	1	3	3	103	9*	3.02×10^4
PI	1	3	3	126	14*	4.04×10^4
RQI	1.5	3	3	187	19	3.24×10^4
PI	1.5	3	3	192	32	3.73×10^4

*tol and upscatter tol = 1×10^{-5}, k tol = 1×10^{-3}

The 3-D benchmark problem shows that for at least some problems, preconditioned RQI converges more quickly in all senses than preconditioned PI if a sufficient amount of preconditioning is used to get RQI to converge. It is pertinent that this is true in the most interesting problem shown so far. However, unpreconditioned PI is still superior in terms of timing.

4.3 PWR 900

Finally, the preconditioner using both PI and RQI was applied to a very challenging problem–a detailed PWR. This problem had 44-groups; $578 \times 578 \times 700$ mesh elements (233,858,800 cells)

broken over 112 × 112 × 10 partitions (12,544 blocks); used a P_0 scattering expansion; used an S_{12} angular quadrature, with S_2 in the preconditioner; and was restricted to a V-cycle depth of 2. The total number of unknowns was thus 1.73 trillion. Based on PWR calculations done previously by Evans and Davidson [6], k is approximately 1.27. The convergence tolerance was 1×10^{-3} and the k tolerance = 1×10^{-3}. This was solved on Titan.

The results using 11 energy sets are given in Table V and show that RQI was much faster and required far fewer Krylov and eigenvalue iterations than PI for this problem. In fact, The preconditioned PI problems did not finish before the wall time limit was reached (RQI was only run with preconditioning since other tests indicated it would not converge within a reasonable time without it). This is the first case where preconditioned RQI was better on all counts than PI with or without preconditioning. RQI with MGE was more than 10 times faster than PI.

Table V. PWR-900: PI and RQI with and without preconditioning, 11 energy sets

Method	Precond	N Eigen	N Krylov	k	Time (m)
RQI	w1r2v2	5	70	$1.268 \pm 1.24 \times 10^{-3}$	54.8
RQI	w1r3v3	6	76	$1.269 \pm 1.12 \times 10^{-3}$	330.4†
PI	none	149	5,602	$1.276 \pm 1.85 \times 10^{-3}$	612.2
PI	w1r2v2	86	946	$1.275 \pm 1.43 \times 10^{-3}$	720*
PI	w1r3v3	11	111	$1.270 \pm 5.09 \times 10^{-2}$	480*,†

*did not converge within walltime limit (8 or 12 hours)
†used S_{12} in MGE preconditioner; full V-cycle depth; tolerance = 1×10^{-4}, k tolerance = 1×10^{-3}; 102 x-blocks, 100 y-blocks, and 10 z-blocks

To investigate how the full system performs for a real problem, a strong scaling study using RQI with MGE using $w1r2v2$ and 1, 4, 11, and 22 sets was done. The results are given in Table VI where $t_{perfect}$ = (1 set solve time / # energy sets) and efficiency = ($t_{perfect}$ / t_{actual}). A strong scaling study with MGE has been published before, but in that case the V-cycle depth was not fixed. This meant that increasing energy sets decreased V-cycle depth such that the preconditioner did less work with more sets. In this study the amount of work done by the preconditioner does not vary with energy sets.

Table VI. PWR-900: RQI strong scaling with $w1r2v2$ preconditioning

Sets	Domains	Time (m)	$t_{perfect}$	Efficiency
1	12,544	407.8	407.8	1.000
4	50,176	123.4	102.0	0.826
11	137,984	54.8	37.1	0.676
22	275,968	39.6	18.5	0.468

The scaling compares quite well to previous scaling studies for Denovo. A fixed source (i.e. MG Krylov only) problem with a similar mesh and 44 groups scaled from 4,320 domains to 190,080 domains with an efficiency of 0.64 [8]. That this problem performed similarly shows that

adding RQI and the MGE preconditioner as solvers does not degrade the strong scaling achieved using the MG Krylov solver only. It is promising that the new solver system does not degrade scaling and that the calculation for which RQI is decisively faster than PI is the one for which this work was designed.

5 CONCLUSIONS

The goal of this research was to accelerate transport calculations with methods that can take full advantage of modern leadership-class computers, facilitating the design of better nuclear systems. Three complimentary methods were used together to demonstrate improvement over traditional methods: the MG Krylov solver, RQI, and the MGE preconditioner. The multigroup Krylov solver converges more quickly than GS and enables energy decomposition such that Denovo can scale to hundreds of thousands of cores. The new multigrid in energy preconditioner reduces iteration count for many problem types and takes advantage of the new energy decomposition such that it can scale very efficiently. These two tools are useful on their own, but together they allow the Rayleigh quotient iteration eigenvalue solver to work. These ideas have never before been used together in this way.

The real motivation of this work was to add RQI, which should converge in fewer iterations than PI for large and challenging problems. RQI creates shifted systems that would not be tractable without the MG Krylov solver. It also creates ill-conditioned matrices that cannot converge without the multigrid in energy preconditioner. Using these methods, RQI converged in fewer iterations and in less time than all PI calculations for a full PWR core. It also scales reasonably well out to 275,968 cores.

The methods added in this research accelerated Denovo in multiple ways. This acceleration helps enable the solution of today's "grand challenge" problems. It is hoped that improved methods will lead to improved reactor designs and systems, and that the frontier of computational challenges will be moved forward.

6 ACKNOWLEDGEMENTS

This research used resources of the Oak Ridge Leadership Computing Facility at the Oak Ridge National Laboratory, which is supported by the Office of Science of the U.S. Department of Energy under Contract No. DE-AC05-00OR22725. Additional thanks to the Rickover Fellowship Program in Nuclear Engineering sponsored by Naval Reactors Division of the U.S. Department of Energy. This fellowship sponsored the work from which this work is derived.

7 REFERENCES

[1] Oak Ridge Leadership Computing Facility, "Introducing Titan, Advancing the Era of Accelerated Computing," `http://www.olcf.ornl.gov/titan/`, 2013.

[2] T. M. Evans, A. S. Stafford, R. N. Slaybaugh, and K. T. Clarno, "Denovo: A New Three-Dimensional Parallel Discrete Ordinates Code in SCALE," *Nuclear Technology*, **171**, *2*, pp. 171–200 (2010).

[3] G. G. Davidson et al., "Massively Parallel, Three-Dimensional Transport Solutions for the k-Eigenvalue Problem," *Nuclear Science and Engineering*, **177**, pp. 111–125 (2014).

[4] R. N. Slaybaugh, T. M. Evans, G. G. Davidson, and P. P. H. Wilson, "Rayleigh Quotient Iteration in 3D, Deterministic Neutron Transport," *PHYSOR 2012 Advances in Reactor Physics Linking Research, Industry, and Education*, Knoxville, TN, 2012, American Nuclear Society.

[5] R. N. Slaybaugh, T. M. Evans, G. G. Davidson, and P. P. H. Wilson, "Multigrid in energy preconditioner for Krylov solvers," *Journal of Computational Physics*, **242**, pp. 405–419 (2013).

[6] T. Evans and G. Davidson, "Technical Note, Subject: Parallel Energy Decomposition in Denovo (Rev. 1)," Oak Ridge National Laboratory (2010).

[7] M. L. Adams and E. W. Larsen, "Fast Iterative Methods for Discrete-Ordinates Particle Transport Calculations," *Progress in Nuclear Energy*, **40**, *1*, pp. 3–159 (2002).

[8] R. N. Slaybaugh, *Acceleration Methods for Massively Parallel Deterministic Transport, a Ph.D. Dissertation*, University of Wisconsin, Madison, WI (2011).

[9] M. A. Heroux et al., "An Overview of the Trilinos Project," *ACM Trans. Math. Softw.*, **31**, *3*, pp. 397–423 (2005).

[10] T. Bailey et al., "Validation of Full-Domain Massively Parallel Transport Sweep Algorithms," *2014 American Nuclear Society Winter Meeting*, Anaheim, CA, 2014, American Nuclear Society.

[11] L. N. Trefethen and D. Bau, III, *Numerical Linear Algebra*, SIAM, Philadelphia, PA (1997).

[12] G. W. Stewart, *Matrix Algorithms Volume II: Eigensystems*, SIAM, Philadelphia, PA, first edition, (2001).

[13] B. N. Parlett, "Rayleigh Quotient Iteration and Some Generalizations for Nonnormal Matrices," *Mathematics of Computation*, **28**, *127*, pp. 679–693 (1974).

[14] G. Peters and J. H. Wilkinson, "Inverse Iteration, Ill-Conditioned Equations and Newton's Method," *Society for Industrial and Applied Mathematics Review*, **21**, *3*, pp. 339–360 (1979).

[15] C. C. Paige, M. Rozložník, and Z. Strakoš, "Modified Gram-Schmidt (MGS), Least Squares, and Backward Stability of MGS-GMRES," *SIAM Journal on Matrix Analysis and Applications*, **28**, *1*, pp. 264–284 (2006).

[16] W. L. Briggs, V. E. Henson, and S. F. McCormick, *A Multigrid Tutorial*, SIAM, Philadelphia, PA, second edition, (2000).

[17] OECD-NEA, "Benchmark on Deterministic Transport Calculations Without Spatial Homogenisation, A 2-D/3-D MOX Fuel Assembly Benchmark," Nuclear Energy Agency, Organisation for Economic Co-operation and Development (2003).

[18] OECD-NEA, "Benchmark on Deterministic Transport Calculations Without Spatial Homogenisation, MOX Fuel Assembly 3-D Extension Case," 5420, Nuclear Energy Agency, Organisation for Economic Co-operation and Development (2005).

[19] S. P. Hamilton, *Numerical Solution of the k-Eigenvalue Problem, a Ph.D. Dissertation*, Emory University, Atlanta, Goergia (2010).